

Controlling Familywise Error When Rejecting at Most One Null Hypothesis Each From a Sequence of Sub-Families of Null Hypotheses

Geoffrey I. Webb
Monash University
`geoff.webb@monash.edu`

Mark van der Laan
University of California, Berkeley
`laan@berkeley.edu`

February 15, 2017

Abstract

We present a procedure for controlling FWER when sequentially considering successive subfamilies of null hypotheses and rejecting at most one from each subfamily. Our procedure differs from previous procedures for controlling FWER by adjusting the critical values that are applied in subsequent rejection decisions by subtracting from the global significance level α quantities based on the p-values of rejected null hypotheses and the numbers of null hypotheses considered.

Keywords: FWER, Sequential Hypothesis Testing, Stepwise Model Selection

1 Introduction

We present a procedure for strictly controlling the Familywise Error Rate when rejecting a single null hypothesis from each subfamily in a sequence of subfamilies of null hypotheses, where each rejection decision is made without knowledge of subsequent subfamilies.

Our procedure is a more powerful variant of a procedure presented by Webb and Petitjean (2016). These procedures differ in form from previous multiple

testing procedures by adjusting the critical value applied to subsequent subfamilies based on the observed values of test statistics for null hypotheses in prior subfamilies.

We identify the assumptions of the procedure, use Monte Carlo simulations to elucidate properties of the procedure under differing scenarios when the assumptions are satisfied, and provide analytical and Monte Carlo simulation results to demonstrate scenarios under which FWER is not controlled when the assumptions are violated.

1.1 Set-up

Let \mathbf{O}_n be a random variable with probability distribution P_0^n . Suppose we observe a realization of this random variable representing our observed data. Let \mathcal{G}_t , $t = 1, \dots, T$, be an ordered sequence of subfamilies of null hypotheses, where $\mathcal{G}_t = \{H_0(t, j) : j = 1, \dots, m_t\}$ consists of m_t null hypotheses about the data distribution P_0^n . Let $T_n(t, j)$ be a test-statistic for null hypothesis $H_0(t, j)$, $j = 1, \dots, m_t$, $t = 1, \dots, T$. Let $\mathcal{N}^t = \{j : H_0(t, j) \text{ is true}\}$ be the set of true null hypotheses in \mathcal{G}_t , and let $\mathcal{F}^t = \{j : H_0(t, j) \text{ is false}\}$ be the set of false null hypotheses in \mathcal{G}_t . Let $\mathcal{N} = \{(t, j) : H_0(t, j) \text{ is true}\}$ and $\mathcal{F} = \{(t, j) : H_0(t, j) \text{ is false}\}$ be the sets of true and false null hypotheses among all null hypotheses.

P-values: Let $P_n(t, j)$ be a p-value implied by $T_n(t, j)$. It is assumed that if $H_0(t, j)$ is true, then $\Pr(P_n(t, j) \leq \alpha) \leq \alpha$ for all $\alpha \in [0, 1]$. In other words, the p -value, which is just a transformation of the test-statistic, satisfies its key property. For simplicity, we assume this to be true for the finite sample $O^{\mathcal{N}}$. As a result, our theorem establishes exact control of the family wise error, but in the often more realistic case that the null distributions of the test statistics are only known asymptotically so that $\lim_{n \rightarrow \infty} \Pr(P_n(t, j) \leq \alpha) \leq \alpha$, our results will provide asymptotic control of the family wise error.

Let $\mathbf{P}_n = (P_n(t, j) : t, j)$ be the vector of p -values, and let $\mathbf{P}_n^{\mathcal{N}} = (P_n(t, j) : (t, j) \in \mathcal{N})$ and $\mathbf{P}_n^{\mathcal{F}} = (P_n(t, j) : (t, j) \in \mathcal{F})$ the vector of p -values for the true null hypotheses and false null hypotheses, respectively. Let $P_n(t) = \min_j P_n(t, j)$ be the minimum p -value for family \mathcal{G}_t and let $J_n(t) = \arg \min_j P_n(t, j)$ identify the null-hypothesis with the minimal p -value. Thus $P_n(t, J_n(t)) = \min_j P_n(t, j)$. We also define $P_n^{\mathcal{N}}(t) = \min_{\{j : (t, j) \in \mathcal{N}\}} P_n(t, j)$ and $P_n^{\mathcal{F}}(t) = \min_{\{j : (t, j) \in \mathcal{F}\}} P_n(t, j)$ as the minimum of the p -values over the set of true and false null hypotheses in family \mathcal{G}_t , respectively. More precisely,

$$P_n^{\mathcal{F}}(t) = \begin{cases} \min_{\{j : (t, j) \in \mathcal{F}\}} P_n(t, j) & \text{if } \mathcal{F}^t \neq \emptyset \\ 1.0 & \text{otherwise} \end{cases}$$

$$P_n^{\mathcal{N}}(t) = \begin{cases} \min_{\{j:(t,j) \in \mathcal{N}\}} P_n(t, j) & \text{if } \mathcal{N}^t \neq \emptyset \\ 1.0 & \text{otherwise} \end{cases}$$

Our goal is to define a sequential multiple testing procedure that rejects at most one hypothesis per subfamily \mathcal{G}_t , making the decision as to whether or not to reject without knowledge of subsequent subfamilies and that controls the familywise error over all subfamilies \mathcal{G}_t , $t = 1, \dots, T$ at user supplied level $\alpha \leq 1$.

1.2 Sequential multiple testing procedure for a sequence of families of null hypotheses.

We propose the following sequential multiple testing procedure that results in a set of rejections $\mathcal{R} = \{J_n(l) : l = 1, \dots, K^*\}$.

Multiple Testing Procedure:

- 1: **let** $K = 1$.
- 2: **let** $\alpha(K) = \alpha$.
- 3: **let** $\mathcal{R} = \emptyset$.
- 4: **let** $K^* = 0$.
- 5: **while** $K \leq T$ and $m_K P_n(K) \leq \alpha(K)$ **do**
- 6: **let** $\mathcal{R} = \mathcal{R} \cup (K, J_n(K))$.
- 7: **let** $\alpha(K + 1) = \alpha(K) - (m_K - 1)P_n(K)$.
- 8: **let** $K = K + 1$.
- 9: **end while**
- 10: **let** $K^* = K - 1$.

This procedure differs from that of Webb and Petitjean (2016) at line 7 where their procedure has instead **let** $\alpha(K + 1) = \alpha(K) - m_K P_n(K)$. By subtracting a smaller quantity from each successive $\alpha(K)$ our procedure is guaranteed to be uniformly more powerful. Hence, our proof also provides a proof of correctness for this prior procedure.

1.3 Theorem establishing family wise error control

The following theorem proves that for each realization of the p -values $\mathbf{P}_n^{\mathcal{F}}$ of the false null hypotheses, the conditional probability of rejecting a true null hypothesis is no greater than α . Of course, this implies, in particular, that the marginal probability on any rejection of a true null is no greater than α . The key assumption this theorem relies upon is that the p -values of the true nulls are independent of the p -values of the false nulls.

Theorem 1 Assume that $\mathbf{P}_n^{\mathcal{N}} = (P_n(t, j) : (t, j) \in \mathcal{N})$ is independent of $\mathbf{P}_n^{\mathcal{F}} = (P_n(t, j) : (t, j) \in \mathcal{F})$. Specifically, assume that for all possible realizations of $\mathbf{P}_n^{\mathcal{F}}$, $Pr(P_n(t, j) \leq \alpha) \leq \alpha$ for all $P_n(t, j) : (t, j) \in \mathcal{N}$ and all $\alpha \in [0, 1]$. Then,

$$Pr(\mathcal{R} \cap \mathcal{N} = \emptyset \mid \mathbf{P}_n^{\mathcal{F}}) \geq 1 - \alpha.$$

Proof: In this proof we condition on $\mathbf{P}_n^{\mathcal{F}}$, so that all probabilities concern the random variable $\mathbf{P}_n^{\mathcal{N}}$.

Scenario I: First, consider the scenario that

$$\sum_{l=1}^{T-1} (m_l - 1)P_n^{\mathcal{F}}(l) + m_T P_n^{\mathcal{F}}(T) \leq \alpha.$$

We note that this implies that all the subfamilies contain at least one false null hypothesis. The probability of a false rejection at the i -th subfamily, $i \in [1, T]$, is the probability that $P_n^{\mathcal{N}}(i) \leq P_n^{\mathcal{F}}(i)$, which is no greater than $(m_i - 1)P_n^{\mathcal{F}}(i)$, where we use that $P_n^{\mathcal{N}}$ is a minimum over maximally $m_i - 1$ true null hypotheses.

The union from $i = 1, \dots, T$ represents the event that we have a false rejection. This proves that the probability of a false rejection is no greater than $\sum_{l=1}^T (m_l - 1)P_n^{\mathcal{F}}(l) < \sum_{l=1}^{T-1} (m_l - 1)P_n^{\mathcal{F}}(l) + m_T P_n^{\mathcal{F}}(T) \leq \alpha$.

Scenario II: The only alternative to Scenario I is the scenario that there exists a first $j \in \{1, \dots, T\}$ such that $m_j P_n^{\mathcal{F}}(j) > \alpha - \sum_{l=1}^{j-1} (m_l - 1)P_n^{\mathcal{F}}(l)$, and thus, for $i = 1, \dots, j-1$, we have $m_i P_n^{\mathcal{F}}(i) \leq \alpha - \sum_{l=1}^{i-1} (m_l - 1)P_n^{\mathcal{F}}(l)$. We note that this implies that the i -th subfamily has at least one false null hypothesis, $i = 1, \dots, j-1$, and that the probability of a false rejection of a true null hypothesis in $\mathcal{G}_1, \dots, \mathcal{G}_{j-1}$ is no greater than $\sum_{l=1}^{j-1} (m_l - 1)P_n^{\mathcal{F}}(l)$.

If there has been no false rejection in $\mathcal{G}_1, \dots, \mathcal{G}_{j-1}$ this implies that for all $i = 1, \dots, j-1$, $P_n(i) = P_n^{\mathcal{F}}(i)$ as otherwise there would have been a false rejection of the true null hypothesis corresponding to $P_n(i)$.

In this scenario, the procedure rejects $J_n(j)$ if and only if $m_j P_n(j) \leq \alpha - \sum_{l=1}^{j-1} (m_l - 1)P_n^{\mathcal{F}}(l)$ and hence the probability of a first false rejection at $\mathcal{G}_j \leq \alpha - \sum_{l=1}^{j-1} (m_l - 1)P_n^{\mathcal{F}}(l)$.

The probability of the union of the two events of a rejection in $\mathcal{G}_1, \dots, \mathcal{G}_{j-1}$ and of a rejection in \mathcal{G}_j but no rejection in $\mathcal{G}_1, \dots, \mathcal{G}_{j-1}$ is thus no greater than $\sum_{l=1}^{j-1} (m_l - 1)P_n^{\mathcal{F}}(l) + \alpha - \sum_{l=1}^{j-1} (m_l - 1)P_n^{\mathcal{F}}(l) = \alpha$.

□

2 Discussion

2.1 Relationship to other approaches for controlling FWER

The standard fixed sequence hypothesis test procedure (Maurer et al., 1995; Hsu and Berger, 1999), where all of a fixed sequence of null hypotheses are tested at level α , is a special case of our procedure where all subfamilies are of size 1 (all $m_i = 1$).

Our procedure follows a fundamentally different strategy to gatekeeping procedures based on Bonferroni adjustments (Bauer et al., 1998; Westfall and Krishen, 2001; Dmitrienko et al., 2003; Chen et al., 2005). Gatekeeping procedures add the $\alpha(i)$ for rejected null hypotheses to the $\alpha(i)$ of subsequent hypotheses. In contrast, our procedure subtracts from subsequent $\alpha(i)$ some portion of the previous $\alpha(i)$, based on the observed p-value of the rejected null hypotheses.

The approach also differs fundamentally from selective inference (Taylor and Tibshirani, 2015). First, our procedure controls FWER, while selective inference controls FDR. Second, unlike our procedure, selective inference, does not use an explicit sequential order over subfamilies of null hypotheses. Third, also unlike our procedure, selective inference rejects null hypotheses in order of ascending p-value until a function over the p-values of the null hypotheses exceed a threshold.

2.2 Monte Carlo experiments

To elucidate the statistical power of the technique, we conducted Monte Carlo simulations. In all the following simulations we use $\alpha = 0.05$.

In the first simulation we generated sets of null hypotheses, which were randomly assigned to be either true or false and were randomly assigned simulated p-values. These simulations were governed by three parameters — *subfamily-Size*: the size of each subfamily; *pTrue*: the probability that a null hypothesis should be designated to be true; and *maxFalsePVal*: the maximum simulated p-value to be assigned to a false null hypothesis.

The following procedure was used for this simulation.

Monte Carlo simulation procedure

```
flag ← false  
 $\mathcal{R} \leftarrow \emptyset$   
sumP ← 0.0  
 $i \leftarrow 0$ 
```

```

while  $flag = \text{false}$  do
   $i \leftarrow i + 1$ 
  Generate  $\mathcal{G}_i$ 
  if  $sumP + |\mathcal{G}_i| \cdot p_i^{\min} \leq \alpha$  then
     $\mathcal{R} \leftarrow \mathcal{R} \cup \{h_i^{\min}\}$ 
     $sumP \leftarrow sumP + |\mathcal{G}_i| \cdot p_i^{\min}$ 
  else
     $flag \leftarrow \text{true}$ 
  end if
end while

```

To generate each \mathcal{G}_i , *subfamilySize* simulated null hypotheses were generated. Each was designated as either true or false, with probability $pTrue$ of being designated true. Each true null hypothesis was assigned a simulated p-value drawn uniformly at random from $[0.0, 1.0]$ and each false null hypothesis was assigned a simulated p-value drawn uniformly at random from $[0.0, 0.1]$. Having lower p-values for false null hypotheses simulates the use of a test statistic that is useful for discriminating between true and false null hypotheses.

$pTrue$ was varied from 0.1 to 1.0 in steps of 0.1 and *subfamilySize* was set to each of the values 1, 10, 100 and 1,000, creating a total of 40 treatments. 1,000,000 Monte Carlo simulations were conducted for each treatment and the FWER and average number of true discoveries per simulation determined.

Figure 1 presents a surface chart showing the effect on FWER as the relative frequency of true to false null hypotheses is increased and as the subfamily size varies. When $pTrue$ is 1.0 and FWER is determined by whether a null hypothesis is rejected for the first subfamily or not, the probability of FWER is strictly controlled by the equivalent of a Bonferroni correction for the first subfamily. FWER falls as the proportion of false null hypotheses rises because the multiple test correction is allowing for the possibility that they are all true.

Increasing subfamilysize also decreases FWER because the multiple test correction allows for the worst case where the rejection regions of all null hypotheses are disjoint whereas in this simulation all null hypotheses are independent of one another.

This simulation demonstrates the power of our procedure when its assumptions are satisfied, and show that it is most powerful when the ratio of false to true hypotheses is highest and subfamilysize is smallest.

We next demonstrate a scenario where violating the requirement that True and False null hypotheses be independent results in a failure to control FWER.

In this scenario we have one false null hypothesis, A and two true null




Figure 1: The FWER of SMT as the relative frequency of true to false null hypotheses is increased and the relative p-values of false relative to true null hypotheses is decreased

hypotheses, B and C . The experimental outcome on which A and B are based is the result of tossing an unbiased coin 17 times. The experimental outcome on which C is based is the result of tossing another coin 13 times. Both coins are unbiased, $Pr(heads) = 0.5$. We choose 17 for the first experiment because it is the smallest number of tosses that has an outcome for a test for $Pr(heads) = 0.5$ that is close to 0.025, and 13 for the second because it is the smallest number of tosses that has an outcome for a test for $Pr(heads) = 0.5$ that is close to 0.05. $A = Pr(heads) \leq 0.1$, $B = Pr(heads) \geq 0.5$ and $C = Pr(heads) \neq 0.5$.

$\mathcal{G}_1 = \{A, B\}$ and $\mathcal{G}_2 = \{C\}$.

We proceed to \mathcal{G}_2 if either A or B is rejected.

A and B are tested at $\alpha/2 = 0.025$.

There are 17 coin tosses and the rejection region for B is 4 or fewer heads. The probability of this outcome is 0.0245.

The rejection region for A is 5 or more heads.

$Pr(17 \text{ heads}) = 7.6294E-06$; $Pr(16 \text{ heads}) = 0.0001$; $Pr(15 \text{ heads}) = 0.0010$; $Pr(14 \text{ heads}) = 0.0052$; $Pr(13 \text{ heads}) = 0.0182, \dots$ The respective p-values for A are $1.00E-17$, $1.54E-15$, $1.117E-13$, $5.0689E-12$, $1.6122E-10$, $3.8152E-09$, $6.9586E-08$, $9.9978E-07$, $1.1464E-05$, 0.0001 , 0.0008 , 0.00467

and 0.0221, meaning A will be rejected if there are 5 or more heads and the adjusted alpha for C will be respectively $0.05 - 1E - 17$ to $0.05 - 0.0221$. If C were a maximally powerful true null hypothesis then the probability of it being rejected would be $7.6294E - 06 \times (\alpha - 1E - 17) + 0.0001 \times (\alpha - 1.54E - 15) + \dots + 0.047 \times 0.0221 = 0.047$. Adding this to the probability of false rejection of B gives a FWER of 0.0715.

However, as we are using coin tosses with a finite number of outcomes, C is not maximally powerful. A Monte Carlo simulation of 1,000,000 repetitions of this scenario yielded a FWER of 0.0647 demonstrating again that violation of the requirement that the true and false null hypotheses be independent of one another can lead to failure to control familywise error.

3 Conclusion

We have presented a novel procedure for controlling familywise error in a sequential testing scenario where at most one null hypothesis is to be rejected from each of a series of subfamilies of null hypotheses. We have shown that this procedure requires only the assumption that the p-values for the true and false null hypotheses are independent of one another. This assumption is realistic in the context of stepwise model selection for which the procedure was developed.

The procedure uses a novel mechanism of adjusting subsequent critical values by quantities based on the observed p-values of null hypotheses that are rejected. It remains a promising avenue for future research to investigate whether this strategy is more broadly applicable in other sequential testing scenarios.

Acknowledgments

This research has been supported by the Australian Research Council under grant DP140100087.

References

- P. Bauer, J. Rhmel, W. Maurer, and L. Hothorn. Testing strategies in multi-dose experiments including active control. *Statistics in Medicine*, 17(18): 2133–2146, 1998. ISSN 1097-0258.

- Xun Chen, Xiaohui Luo, and Tom Capizzi. The application of enhanced parallel gatekeeping strategies. *Statistics in Medicine*, 24(9):1385–1397, 2005. ISSN 1097-0258.
- Alexei Dmitrienko, Walter W. Offen, and Peter H. Westfall. Gatekeeping strategies for clinical trials that do not require all primary effects to be significant. *Statistics in Medicine*, 22(15):2387–2400, 2003. ISSN 1097-0258.
- Jason C. Hsu and Roger L. Berger. Stepwise confidence intervals without multiplicity adjustment for doseresponse and toxicity studies. *Journal of the American Statistical Association*, 94(446):468–482, 1999.
- W Maurer, LA Hothorn, and W Lehmacher. Multiple comparisons in drug clinical trials and preclinical assays: a-priori ordered hypotheses. *Biometrie in der chemisch-pharmazeutischen Industrie*, 6:3–18, 1995.
- Jonathan Taylor and Robert J. Tibshirani. Statistical learning and selective inference. *Proceedings of the National Academy of Sciences*, 112(25):7629–7634, 2015.
- Geoffrey I Webb and Francois Petitjean. A multiple test correction for streams and cascades of statistical hypothesis tests. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD-16*, pages 1255–1264. ACM Press, 2016.
- Peter H. Westfall and Alok Krishen. Optimally weighted, fixed sequence and gatekeeper multiple testing procedures. *Journal of Statistical Planning and Inference*, 99(1):25 – 40, 2001. ISSN 0378-3758.